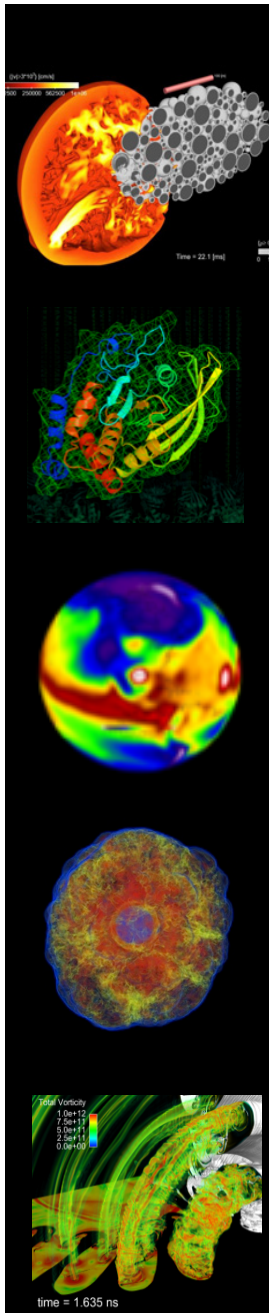


# Vendor Roadmap Presentation Guidance

James Lujan, James Laros



Los Alamos National Laboratory, an affirmative action/equal opportunity employer, is operated by Los Alamos National Security, LLC, for the National Nuclear Security Administration of the U.S. Department of Energy under contract DE-AC52-06NA25396. LA-UR-18-26721. Approved for public release; distribution is unlimited.

Sandia National Laboratories is a multi-mission laboratory managed and operated by National Technology & Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525. SAND2018-7955 PE



# Why are we here?

---

- This is a technical meeting!
- We have our Subject Matter Experts present to hear about your technologies
- Please defer any non-technical questions (e.g. business, procurement, terms and conditions) to the ACES leadership at another time.

# What/Who is ACES?

---

- ACES (*New Mexico **A**lliance for **C**omputing at **E**xtr~~e~~m~~e~~ **S**cale*) is a collaboration between Los Alamos National Laboratory (LANL) and Sandia National Laboratories (SNL)
- The previous Crossroads RFP was released under APEX (***A**lliance for **a**pplication **P**erformance at **E**Xtr~~e~~m~~e~~ **s**cale*)
  - APEX = ACES + NERSC
- NERSC no longer involved therefore RFP will be released by ACES and reflect ACES requirements only
- Single platform procurement (Crossroads)
- Target hardware delivery by end of FY21
- Target acceptance by end of CY21

# High-Level Design Philosophy

---

- The goal of the Crossroads platform procurement is **Efficiency**.
- Efficiency will be evaluated in the areas of:
  - **Performance efficiency**
  - **Workflow efficiency**
  - **Porting efficiency**
- Performance efficiency is defined as the achieved performance of the application once ported to the proposed platform.
- Workflow efficiency is defined as the efficiency that a complete NNSA workflow executes on the proposed platform.
- Porting efficiency is defined as the ease in which NNSA mission codes can be ported to execute on the proposed architecture. Minimal change to the existing code base is of high value.
- When evaluating proposals efficiency in all three stated areas will be considered together.



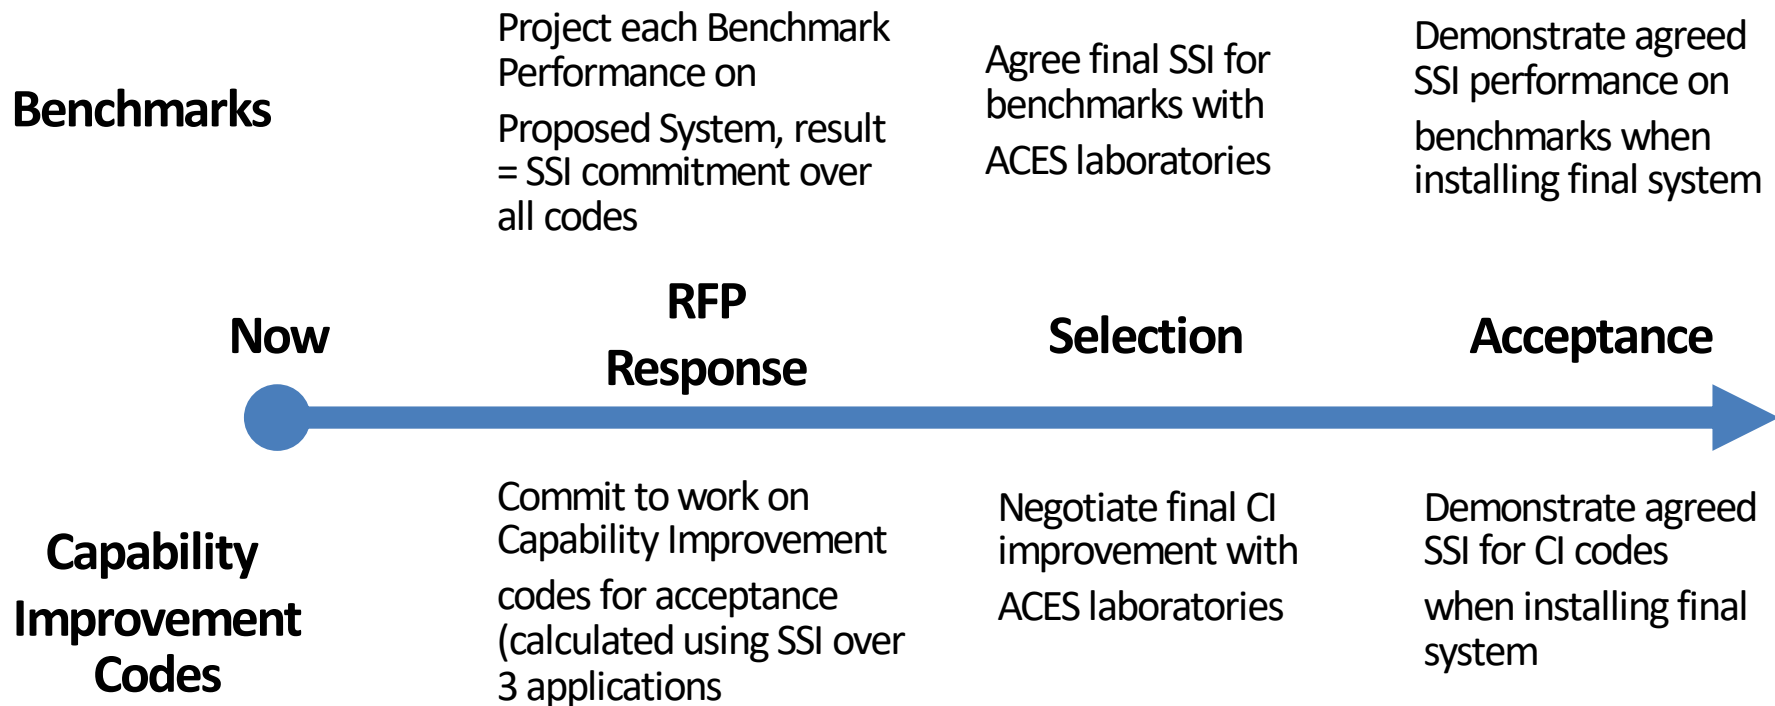
# 'Application Performance'

---

- Crossroads Benchmarks
  - SNAP, HPCG, PENNANT, MiniPIC, UMT, VPIC, Branson
- Capability Improvement Applications (ASC Simulations Codes)
  - Mercury (LLNL), PartiSn (LANL), SPARC (SNL)
- Microbenchmarks
  - DGEMM, IOR, MDTest, MPI, Stream
- → Metrics used for system acceptance (e.g. SSI)

# Benchmarks and Capability Improvement Applications

---



# Benchmarks

---

- Benchmarks can be downloaded now
  - SNAP, HPCG, PENNANT, MiniPIC, UMT, VPIC, Branson
  - Unrestricted applications
  - Mostly MPI + OpenMP, Mix of Fortran and C/C++
- At time of response, commit to a system SSI over all benchmarks
  - Baseline SSI (mostly unmodified codes)
  - Optimized SSI (can optimize to show off the system/software)

# Capability Improvement

---

- Capability Improvement codes
  - Mercury (LLNL), PartiSn (LANL), SPARC (SNL)
  - Export Controlled codes (OD999 and ITAR level)
  - Mostly MPI, some OpenMP/Kokkos
  - Mix of Fortran, C/C++
- At time of response commit to work on these codes for acceptance, agree on final SSI over all three codes during negotiation with ACES (once selected)

# 'Application Transition Support' - Section 3.7 of Tech Specs

---

- Collaboration (i.e. a Center of Excellence) between vendor experts in the areas of application porting and performance optimization and the Crossroads application development community.
  - Vendor activities can support:
    - Training
    - Workshop, Hackathon, Discovery (2-3 days at customer site, 8+ codes)
    - Non-export controlled deep dive (6-8 weeks preparation, 3 days at vendor site)
    - Export controlled deep dive (6-8 weeks preparation, 3 days at vendor site using ACES compute resources)
  - Provided from the date of subcontract execution through two (2) years after final acceptance of the systems.
  - NO Predefined metrics (collaboration)
  - Definitely includes Export Controlled codes
  - Vendor staff with a DOE Q-clearance a plus (+), embedded on-site support (++)
- Support structure for the proposed programming environment (e.g. reporting issues, requesting new functionality, escalation paths/priorities available to Crossroads' applications).
  - Support up to two (2) years after final acceptance of the systems.

# Facility, Power & Cooling

---

- Crossroads will be located in the Nicholas C. Metropolis center (SCC) at Los Alamos National Lab
- Estimated facility power and footprint
  - Crossroads
    - 15MW (3-phase 480V)
    - 8000 square feet
- Liquid cooled
  - Is our assumption correct?
  - Warm water or chilled ? Direct or indirect?

# Guiding Questions

---

- Basically we want to understand your roadmap(s) in the timeframe we anticipate taking delivery (FY2021)
- Your roadmap presentations should NOT be limited to these guiding questions
- Tell us where and why our assumptions are wrong!
- What memory technology (technologies)?
  - DDRx?
  - On-package (HBM or other?)
- What does the memory hierarchy look like?
  - Single/Multi-tier?
  - NUMA?
  - Bandwidth and latency characteristics (between levels or NUMA regions)?
  - Capacity?
  - Relative cost and energy trade-offs?
- What is the node architecture?
  - Processor technology?
  - How many cores?
  - Heterogeneous or Homogeneous?
  - Core characteristics
  - Coherency?

# Guiding Questions (cont)

---

- High speed network
  - We want one 😊
  - Technology?
  - Topology ?
  - NIC?
    - Integrated or discrete?
  - Injection bandwidth?
  - Bisection bandwidth?
  - Message injection rate?
    - At what message size(s)?
  - Offload characteristics?
  - Access to memory?

# Guiding Questions (cont)

---

- Software
  - Languages
  - Programming Environments
  - Programming Models
  - Profilers and Debuggers
  - Operating system(s)
  - Advanced Power Measurement and Control
  - RAS and/or System Management
  - Software to aid resiliency
  - Workload (and workflow) management

# Guiding Questions (cont)

---

- What will the file system look like?
  - Integrated into memory hierarchy?
  - Is traditional application driven check point restart still required?
  - How can we optimize for analysis usage models?
- Support for task based programming model(s)?
- What are the advanced resilience mechanisms?
  - Hardware and/or software
- Will you have early test platforms / proxies available that we can explore these issues with?
- How can Crossroads best influence your roadmap?

# Schedule

---

- Draft Tech Spec's released July 2018
- Vendor briefings week of July 23<sup>rd</sup>, 2018  
Albuquerque NM
- RFP released Fall 2018
- 45 day response to RFP
  - Quiet time during response period